arXiv:1401.7043v2 [cs.DM] 20 Sep 2014

Randomized Minmax Regret for Combinatorial Optimization Under Uncertainty*

Andrew Mastin†, Patrick Jaillet‡, Sang Chin§

September 23, 2014

Abstract

The minmax regret problem for combinatorial optimization under uncertainty can be viewed as a zero-sum game played between an optimizing player and an adversary, where the optimizing player selects a solution and the adversary selects costs with the intention of maximizing the regret of the player. The existing minmax regret model considers only deterministic solutions/strategies, and minmax regret versions of most polynomial solvable problems are NP-hard. In this paper, we consider a randomized model where the optimizing player selects a probability distribution (corresponding to a mixed strategy) over solutions and the adversary selects costs with knowledge of the player's distribution, but not its realization. We show that under this randomized model, the minmax regret version of any polynomial solvable combinatorial problem becomes polynomial solvable. This holds true for both the interval and discrete scenario representations of uncertainty. Using the randomized model, we show new proofs of existing approximation algorithms for the deterministic model based on primal-dual approaches. Finally, we prove that minmax regret problems are NP-hard under general convex uncertainty.

1 Introduction

Many optimization applications involve cost coefficients that are not fully known. When distributional information on cost coefficients is available (e.g. from historical data or other estimates), stochastic programming is often an appropriate modeling choice [13, 24]. In other cases, costs may only be known to be contained in intervals (i.e. each cost has a known lower and upper bound), or to be a member of a finite set of scenarios, and one is more interested in worst-case performance. Robust optimization formulations are desirable here as they employ a minmax-type objective and do not require knowledge of cost distributions [9, 18, 21].

In a general robust optimization problem with cost uncertainty, one must select a set of items from some feasible *solution set*, such that item costs are unknown but must be contained in a known

*Research supported in part by NASA ESTOs Advanced Information System Technology (AIST) program under grant number NNX12H81G. Also supported by NSF grant 1029603 and ONR grant N00014-12-1-0033.

†Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; mastin@mit.edu

‡Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science and Operations Research Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; jaillet@mit.edu

§Draper Laboratory, 555 Technology Square, Cambridge, MA 02139; schin@draper.com

uncertainty set. Under the well known *minmax* objective (also referred to as absolute robustness), the goal is to select a solution that gives the best upper bound on objective cost over all possible costs from the uncertainty set [25]. That is, one must select the solution that, when item costs are chosen to maximize the cost of the selected solution, is minimum. Under the *minmax regret* objective (sometimes called the robust deviation model), the goal is instead to select the solution that minimizes the maximum possible regret, defined as the difference between the cost of the selected solution and the optimal solution [23].

A problem under the minmax regret objective can be viewed as a two stage game. In the first stage, the optimizing player selects a deterministic solution. In the second stage, an adversary observes the selected solution and chooses costs from the uncertainty set with the intention of maximizing the player's regret. The goal of the optimizing player is thus to select a solution that least allows the adversary to generate regret. For both interval and discrete scenario representations of cost uncertainty, the minmax regret versions of most polynomial solvable problems are NP-hard [4]. A variation on this model, first suggested by Bertsimas et al. [11] for minmax robust optimization, is to allow the optimizing player to select a probability distribution over solutions and require the adversary to select costs based only on knowledge of the players distribution, but not its realization. In this paper, we show that under this randomized model, the minmax regret version of any polynomial solvable 0-1 integer linear programming problem becomes polynomial solvable. This holds true for both the interval and discrete scenario representations of uncertainty.

Our crucial observation is that the randomized model is the linear programming relaxation of the integer program for the deterministic model. This leads to some useful insights. First, the minmax expected regret in the randomized model is upper bounded by the minmax regret in the deterministic model. Next, the linear program formulation can be used to create an approximation algorithm for the deterministic problem. We show that existing approximation algorithms for deterministic minmax regret problems, which have been proved using combinatorial arguments, can in fact be derived using primal-dual methods [3, 19]. Our analysis here leads to lower bounds on randomized minmax regret with respect to the deterministic minmax regret, effectively stating limits on the power of using randomization.

Given that the randomized model makes the minmax regret problem polynomial solvable for interval uncertainty and discrete scenario uncertainty, it is natural to ask if polynomial solvability remains in the presence of slightly more elaborate uncertainty sets. We show that for general convex uncertainty sets, however, that the mere maximum regret problem (rather than the full minmax regret problem) is NP-hard. The deterministic and randomized minmax regret problems are at least as hard as the maximum regret problem, so these problems become NP-hard under general convex uncertainty.

The paper is structured as follows. In the remainder of this section we review related work; Section 2 introduces notation and definitions. Section 3 presents the analysis for discrete scenario uncertainty, with derivations of optimal strategies for the optimizing player and the adversary, as well as the primal-dual approximation algorithm. Section 4 gives the same results for interval uncertainty. Section 5 demonstrates NP-hardness of minmax regret problems under general convex uncertainty. A conclusion is given in Section 6.

1.1 Related Work

One of the first studies of minmax regret from both an algorithmic and complexity perspective was that of Averbakh [5]. He looked at the minmax regret version of the simple problem of selecting k

items out of n total items where the cost of each item is uncertain, and the goal is to select the set of items with minimum total cost. For interval uncertainty, he derived a polynomial time algorithm based on interchange arguments. He demonstrated that for the discrete scenario representation of uncertainty, however, the minmax regret problem becomes NP-hard, even for the case of only two scenarios. It is interesting to contrast these results with the case of general minmax regret linear programming, which as shown by Averbakh and Lebedev [7], is NP-hard for interval uncertainty but polynomial solvable for discrete scenario uncertainty.

Apart from the item selection problem, most polynomial solvable minmax regret combinatorial problems are NP-hard, both for interval and discrete scenario uncertainty. This is true for the shortest path, minimum spanning tree, assignment, and minimum s-t cut problems [1, 2, 6, 21, 27]. One exception is the minimum cut problem, the minmax regret version of which is polynomial solvable both for interval and discrete scenario uncertainty [2]. The survey paper of Aissi et al. [4] provides a comprehensive summary of results related to both minmax and minmax regret combinatorial problems. For problems that are already NP-complete, most of their minmax regret versions are Σ_2^p-complete (meaning that they are at the second level of the polynomial hierarchy) [17]. To solve minmax regret problems in practice, the book by Kasperski reviews standard mixed integer program (MIP) formulations for both interval and discrete scenario uncertainty [18]. General approximation algorithms are known for both types of uncertainty. Kasperski and Zieliński [19] proved a general 2-approximation algorithm based on midpoint costs under interval uncertainty, and Aissi et al. [3] gave a k-approximation algorithm using average costs under discrete scenario uncertainty, where k is the number of scenarios.

The application of a game theoretic model with mixed strategies to robust optimization problems was introduced by Bertsimas et al. [11]. They focused on the minmax robust model, and their analysis was motivated by adversarial models used for online optimization algorithms. As described by Ben-David et al. [8] (see also Borodin and El-Yaniv [14]), the three types of adversaries are the *oblivious adversary*, the *adaptive online adversary*, and the *adaptive offline adversary*. The adaptive offline adversary is the analog of the conventional deterministic minmax regret problem, while the adaptive online adversary corresponds to our randomized model. The analog of the oblivious adversary, which we do study, is the model where the adversary first selects costs, and the optimizing player then selects the solution after viewing these costs.

For the randomized (corresponding to the adaptive online adversary) minmax problem, Bertsimas et al. [11] showed that if it is possible to optimize over both the solution set and the uncertainty set in polynomial time, then an optimal mixed strategy solution can be calculated in polynomial time, and that the expected cost under the randomized model is no greater than the cost for the deterministic model. This holds despite the fact that solving the minmax version of many polynomial solvable problems is NP-hard for the deterministic case [10]. They also gave lower bounds on the improvement gained from randomization for various uncertainty sets. Our work is similar to theirs, but we focus on the minmax regret objective instead of the minmax objective.

Another line of research that is related to ours is in security applications, where the adversarial model is realistically motivated. Korzhyk, et al. [20] considered assignment-type problems where defensive resources, such as security guards, must be assigned to valued targets. They followed a Stackelberg model where the defending player has the power to commit to a mixed strategy; the attacker then observes this mixed strategy (though not the realization) and decides which targets to attack. They used linear programming formulations along with the Birkhoff-von Neumann theorem to find polynomial-sized optimal mixed strategies. It is also worth mentioning the work of Bertismas

et al. on randomized strategies for network interdiction [12].

2 Definitions

We consider a general combinatorial optimization problem where we are given a set of n items $E = \{e_1, e_2, \ldots, e_n\}$ and a set $\mathcal{F}$ of feasible subsets of E. Each item $e \in E$ has a cost $c_e \in \mathbb{R}$. Given the vector $c = (c_1, \ldots, c_n)$, the goal of the optimization problem is to select the feasible subset of items that minimizes the total cost; we refer to this as the *nominal problem*:

$$F^*(c) := \min_{T \in \mathcal{F}} \sum_{e \in T} c_e. \tag{1}$$

Let $x = (x_1, \ldots, x_n)$ be a characteristic vector for some set T, so that $x_e = 1$ if $e \in T$ and $x_e = 0$ otherwise. Also let $\mathcal{X} \subseteq \{0,1\}^n$ denote the set of all characteristic vectors corresponding to feasible sets $T \in \mathcal{F}$. We assume that $\mathcal{X}$ is described in size m (e.g. with m linear inequalities). We can equivalently write the nominal problem with a linear objective function:

$$F^*(c) = \min_{x \in \mathcal{X}} \sum_{e \in E} c_e x_e. \tag{2}$$

Throughout the paper, we will use both set notation and characteristic vectors for ease of presentation.

We will review the conventional regret definitions for the deterministic minmax regret framework, and then present the analogous definitions for our randomized model. For some cost vector $c \in \mathcal{C}$, the deterministic cost of a solution $T \in \mathcal{F}$ is

$$F(T, c) := \sum_{e \in T} c_e. \tag{3}$$

The regret of a solution T under some cost vector c is the difference between the cost of the solution and the optimal cost:

$$R(T, c) := F(T, c) - F^*(c). \tag{4}$$

The *maximum regret problem* for a solution T is

$$R_{\max}(T) := \max_{c \in \mathcal{C}} R(T, c) = \max_{c \in \mathcal{C}} \left(F(T, c) - F^*(c) \right). \tag{5}$$

The *deterministic minmax regret problem* is then

$$Z_{\mathrm{D}} := \min_{T \in \mathcal{F}} R_{\max}(T) = \min_{T \in \mathcal{F}} \max_{c \in \mathcal{C}} (F(T, c) - F^*(c)). \tag{6}$$

In the remainder of the paper, we will frequently abuse the notation $F(\cdot, c)$, $R(\cdot, c)$ and $R_{\max}(\cdot)$ by replacing set arguments with vectors (e.g. $F(x, c)$ in place of $F(T, c)$), but we will follow the convention of using capital letters for sets and lowercase letters for vectors.

We now move to the randomized framework, where the optimizing player selects a distribution over solutions and the adversary selects a distribution over costs. Starting with the optimizing player, for some set $T \in \mathcal{F}$, let y_T denote the probability that the optimizing player selects set T.

Let $y = (y_T)_{T \in \mathcal{F}}$ be the vector of length $|\mathcal{F}|$ specifying the set selection distribution; we will refer to y simply as a solution. Define the feasible region for y as

$$\mathcal{Y} := \{y | y \geq \mathbf{0}, \mathbf{1}^\top y = 1\}, \tag{7}$$

where the notation $\mathbf{0}$ and $\mathbf{1}$ indicates a full vector of zeros and ones, respectively. We similarly define a distribution over costs for the adversary. The set $\mathcal{C}$ may in general be infinite, but we will only consider strategies with finite support; for now we will assume that such strategies are sufficient. Thus consider a finite set $\mathcal{C}_f \subseteq \mathcal{C}$, and for some $c \in \mathcal{C}_f$, let w_c denote the probability that the adversary selects costs c. Then let $w = (w_c)_{c \in \mathcal{C}_f}$ and define the feasible region

$$\mathcal{W} := \{w | w \geq \mathbf{0}, \mathbf{1}^\top w = 1\}. \tag{8}$$

The expected regret under y and w is simply

$$\overline{R}(y, w) := \sum_{T \in \mathcal{F}} \sum_{c \in \mathcal{C}_f} y_T w_c R(T, c) = \sum_{T \in \mathcal{F}} \sum_{c \in \mathcal{C}_f} y_T w_c (F(T, c) - F^*(c)). \tag{9}$$

For a given y, the *maximum expected regret problem* is

$$\begin{aligned} \overline{R}_{\max}(y) &:= \max_{w \in \mathcal{W}} \sum_{c \in \mathcal{C}_f} w_c \sum_{T \in \mathcal{F}} y_T R(T, c) \\ &= \max_{c \in \mathcal{C}_f} \sum_{T \in \mathcal{F}} y_T R(T, c). \end{aligned} \tag{10}$$

The above equality follows using the standard observation used in game theory: the optimization of $w \in \mathcal{W}$ is maximization of the function $G(y, c) = \sum_{T \in \mathcal{F}} y_T R(T, c)$ over the convex hull of $\mathcal{C}_f$, which is equivalent to optimizing over $\mathcal{C}_f$ itself. The minmax expected regret problem, which we refer to as the *randomized minmax regret problem*, is

$$Z_{\mathrm{R}} := \min_{y \in \mathcal{Y}} \overline{R}_{\max}(y) = \min_{y \in \mathcal{Y}} \max_{c \in \mathcal{C}} \left(\sum_{T \in \mathcal{F}} y_T (F(T, c) - F^*(c)) \right), \tag{11}$$

where we have replaced $\mathcal{C}_f$ with $\mathcal{C}$ under the assumption that $\mathcal{C}_f$ contains the maximizing cost vector.

The above minmax expected regret problem is the problem faced by the optimizing player; the adversary, however, is interested in solving the maxmin expected regret problem, defined as follows. First, the *minimum expected regret problem* for a given w is

$$\begin{aligned} \overline{R}_{\min}(w) &:= \min_{y \in \mathcal{Y}} \sum_{T \in \mathcal{F}} y_T \sum_{c \in \mathcal{C}_f} w_c R(T, c) \\ &= \min_{T \in \mathcal{F}} \sum_{c \in \mathcal{C}_f} w_c R(T, c), \end{aligned} \tag{12}$$

where we have once again used the fact that optimizing over the convex hull of the set of solutions is equivalent to optimizing over the set of solutions. The *adversarial randomized maxmin regret problem* is

$$Z_{\mathrm{AR}} := \max_{w \in \mathcal{W}} \overline{R}_{\min}(w) = \max_{w \in \mathcal{W}} \min_{T \in \mathcal{F}} \left(\sum_{c \in \mathcal{C}_f} w_c (F(T, c) - F^*(c)) \right). \tag{13}$$

It is often the case that the minmax value of the game is equal to the maxmin value; that is, $Z_{\mathrm{R}} = Z_{\mathrm{AR}}$. The Minimax Theorem states that this holds for two-person zero-sum games with a finite number of pure strategies [22]. In the following sections, we will show that this identity holds for discrete scenario uncertainty and interval uncertainty, following from linear programming duality.

3 Discrete Scenario Uncertainty

Under discrete scenario uncertainty, we are given a finite set $\mathcal{S}$ of $|\mathcal{S}| = k$ scenarios. For each $S \in \mathcal{S}$, there exists a cost vector $c^S = (c_e^S)_{e \in E}$. The adversary's mixed strategy is a probability distribution over scenarios, so we are not concerned with complications arising from infinite sets. This section is divided into three parts; we first determine computation of the optimal randomized strategy for the optimizing player, followed by computation of the adversary's optimal strategy. Thereafter, we use the randomized model to devise a primal-dual approximation scheme for the deterministic minmax regret problem. We restate and clarify some notation in the context of discrete scenario uncertainty throughout our development.

3.1 Optimizing Player

We first make some observations regarding the deterministic minmax regret problem that will be helpful in making comparisons with the randomized model. Under discrete scenario uncertainty, the deterministic maximum regret problem is

$$R_{\max}(T) = \max_{S \in \mathcal{S}} R(T, c^S) = \max_{S \in \mathcal{S}} \left(F(T, c^S) - F^*(c^S) \right). \tag{14}$$

The deterministic minmax regret problem is

$$Z_{\mathrm{D}} = \min_{T \in \mathcal{F}} R_{\max}(T) = \min_{T \in \mathcal{F}} \max_{S \in \mathcal{S}} (F(T, c^S) - F^*(c^S)). \tag{15}$$

Lemma 1. *The deterministic minmax regret problem with discrete scenario uncertainty is equivalent to the following integer program.*

$$\begin{aligned} Z_{\mathrm{D}} \;=\; \min \quad & z & (16)\\ \text{s.t.} \quad & \sum_{e \in E} c_e^S x_e - F^*(c^S) \le z, \qquad \forall S \in \mathcal{S}, \\ & x \in \mathcal{X}. \end{aligned}$$

Proof. Slightly abusing the notation for maximum regret, we have with vector notation

$$R_{\max}(x) = \max_{S \in \mathcal{S}} \left(\sum_{e \in E} c_e^S x_e - F^*(c^S) \right). \tag{17}$$

The integer program then follows by definition of the maximum. □

For the randomized model, recall that the optimizing player's distribution over solutions is denoted by $y = (y_T)_{T \in \mathcal{F}}$ and that $\mathcal{Y}$ denotes the set of valid probability distributions. The maximum

expected regret problem is

$$\begin{aligned}\overline{R}_{\max}(y) &= \max_{S\in\mathcal{S}} \sum_{T\in\mathcal{F}} y_T R(T, c^S) \\ &= \max_{S\in\mathcal{S}} \left(\sum_{T\in\mathcal{F}} y_T F(T, c^S) - F^*(c^S) \right). \end{aligned} \tag{18}$$

We define the expected value of a solution for a distribution y and cost vector c^S to simplify notation:

$$\overline{F}(y, c^S) := \sum_{T\in\mathcal{F}} y_T F(T, c^S). \tag{19}$$

The maximum expected regret problem can then be stated as

$$\overline{R}_{\max}(y) = \max_{S\in\mathcal{S}} (\overline{F}(y, c^S) - F^*(c^S)). \tag{20}$$

The randomized minmax regret problem is

$$Z_{\mathrm{R}} = \min_{y\in\mathcal{Y}} \overline{R}_{\max}(y) = \min_{y\in\mathcal{Y}} \max_{S\in\mathcal{S}} (\overline{F}(y, c^S) - F^*(c^S)). \tag{21}$$

To solve the randomized minmax regret problem, it is possible to write a linear program analogous to the above integer program using variables y_T. This would, however, have $|\mathcal{F}|$ variables, which may grow exponentially in n. Instead, we note that for the maximum regret expected regret problem,

$$\begin{aligned}\overline{R}_{\max}(y) &= \max_{S\in\mathcal{S}} \left(\sum_{T\in\mathcal{F}} y_T \sum_{e\in T} c_e^S - F^*(c^S) \right) \\ &= \max_{S\in\mathcal{S}} \left(\sum_{e\in E} c_e^S \sum_{T\in\mathcal{F}: e\in T} y_T - F^*(c^S) \right). \end{aligned} \tag{22}$$

The change in summation order motivates the substitution

$$p_e := \sum_{T\in\mathcal{F}: e\in T} y_T, \quad e \in E. \tag{23}$$

Let $p = (p_1, \ldots, p_n)$; we will refer to this as the marginal probability vector. The substitution is a mapping from $\mathcal{Y}$ to the convex hull of $\mathcal{X}$. The following is the minmax regret analog of an observation made by Bertsimas et al. [11].

Lemma 2. *For discrete scenario uncertainty, the objective value Z_{R} of the randomized minmax regret problem* (11) *is equal to that of the problem*

$$\min_{p\in \mathrm{CH}(\mathcal{X})} \max_{S\in\mathcal{S}} \left(\sum_{e\in E} c_e^S p_e - F^*(c^S) \right), \tag{24}$$

where $\mathrm{CH}(\mathcal{X})$ *denotes the convex hull of* $\mathcal{X}$*.*

Proof. We use the same arguments presented in [11]. By definition of the substitution (23), the vector p must lie in the convex hull of $\mathcal{X}$. Carathéodory's Theorem [15] states that any $p \in \text{CH}(\mathcal{X})$ can be represented by a convex combination of at most $n+1$ points in $\mathcal{X}$, so there exists a surjective mapping from $\mathcal{Y}$ to $\text{CH}(\mathcal{X})$. □

Since we will use the simplified formulation given in Lemma 2 to solve the randomized minmax regret problem, we address the problem of recovering a vector y given a solution p. In the proof of the lemma, we have used Carathéodory's Theorem, which proves existence of such a mapping, but not its construction. To this end, we define for the optimizing player a *mixed strategy encoding* $\mathcal{M} = (X, Y)$ as a set of deterministic solutions $X = \{x^{T_i} \in \mathcal{X},\ i = 1, \ldots, \mu\}$ that should be selected with nonzero probability and the corresponding probabilities $Y = \{y_{T_i} \in [0,1],\ i = 1, \ldots, \mu\}$ that satisfy $\sum_{i=1}^{\mu} y_{T_i} = 1$. Here μ is the support size of the mixed strategy (i.e. the number of deterministic solutions with nonzero probablity). For a given vector p, we are interested in solving the following constraint satisfaction program:

$$\begin{aligned} \min \quad & 0 && (25)\\ \text{s.t.} \quad & \sum_{T \in \mathcal{F}: e \in T} y_T = p_e, \qquad \forall e \in E, \\ & \sum_{T \in \mathcal{F}} y_T = 1, \\ & y \geq \mathbf{0}. \end{aligned}$$

Consider the dual program of (25), which has variables $u = (u_1, \ldots, u_e)$ and w:

$$\begin{aligned} \max \quad & w - \sum_{e \in E} p_e u_e && (26)\\ \text{s.t.} \quad & w - \sum_{e \in T} u_e \leq 0, \qquad \forall T \in \mathcal{F}, && (27)\\ & u, w \text{ free.} \end{aligned}$$

Recall that the region $\mathcal{X}$ is described in size m.

Lemma 3. *For any given $p \in \text{CH}(\mathcal{X})$, a corresponding mixed strategy encoding $\mathcal{M}$ of size polynomial in n can be found via the linear programming formulation* (26) - (27)*. Furthermore, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then $\mathcal{M}$ can be found in time polynomial in n and m.*

Proof. Notice that while the primal program has an exponential number of variables and a linear number of constraints, the opposite holds true for the dual. The primal program is bounded since all objective coefficients are equal to zero, and is feasible due to Carathéodory's Theorem. Therefore the dual program must be feasible and bounded.

To guarantee a polynomial sized solution, note that the separation problem for the constraints (27) is simply the nominal problem with costs u, so the dual program can be solved via the ellipsoid method. If the nominal problem can be solved in polynomial time, then the constraints (27) can be generated in polynomial time, giving a polynomial time solution for the entire dual program (26) - (27).

From a practical perspective, a separation oracle for (27) gives an efficient method for performing row generation with the simplex method. Each row i generated while solving the dual problem gives a solution $x^{T_i} \in \mathcal{X}$, and its dual variable is the corresponding probability y_{T_i}. □

Using Lemma 2, we can now formulate a linear program to solve the randomized minmax regret problem.

$$\begin{aligned} \min \quad & z && (28)\\ \text{s.t.} \quad & \sum_{e \in E} c_e^S p_e - F^*(c^S) \leq z, \qquad \forall S \in \mathcal{S}, && (29)\\ & p \in \mathrm{CH}(\mathcal{X}). && (30) \end{aligned}$$

This leads to the important result that the randomized minmax regret problem is polynomial solvable for any polynomial solvable nominal problem. Also, the minmax expected regret is upper bounded by the minmax regret in the deterministic case.

Theorem 1. *For discrete scenario uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized minmax regret problem $\min_{y \in \mathcal{Y}} \max_{S \in \mathcal{S}} (\overline{F}(y, c^S) - F^*(c^S))$ can be solved in time polynomial in n, m, and k.*

Proof. Since for all $S \in \mathcal{S}$, the value $F^*(c^S)$ is polynomial solvable, each constraint (29) can be enumerated in polynomial time. If we can optimize over $\mathcal{X}$ in polynomial time, then we can separate over $\mathrm{CH}(\mathcal{X})$ in polynomial time via the result of [16]. This gives the separation oracle for (30). □

Corollary 1. *For discrete scenario uncertainty, $Z_{\mathrm{R}} \leq Z_{\mathrm{D}}$.*

Proof. The program (28) - (30) is the linear programming relaxation of (16). □

3.2 Adversary

Moving to the perspective of the adversary under discrete scenario uncertainty, the adversary must select a mixed strategy over scenarios. The finite number of scenarios naturally requires the adversary's distribution to have finite support. Specifically, the adversary selects a distribution over costs $w = (w_S)_{S \in \mathcal{S}}$. The minimum expected regret problem for a given w is

$$\begin{aligned} \overline{R}_{\min}(w) &= \min_{T \in \mathcal{F}} \sum_{S \in \mathcal{S}} w_S R(T, c^S) \\ &= \min_{T \in \mathcal{F}} \sum_{S \in \mathcal{S}} w_S \left(F(T, c^S) - F^*(c^S) \right). \end{aligned} \tag{31}$$

Recall that $\mathcal{W}$ indicates valid probability distributions for w. The adversarial randomized maxmin regret problem is

$$Z_{\mathrm{AR}} = \max_{w \in \mathcal{W}} \overline{R}_{\min}(w) = \max_{w \in \mathcal{W}} \min_{T \in \mathcal{F}} \sum_{S \in \mathcal{S}} w_S \left(F(T, c^S) - F^*(c^S) \right). \tag{32}$$

From the above definition, we formulate a linear program to solve the adversarial randomized maxmin regret problem:

$$\max \quad z \tag{33}$$

$$\text{s.t.} \quad \sum_{S\in\mathcal{S}} w_S(F(T,c^S) - F^*(c^S)) \geq z, \qquad \forall T \in \mathcal{F}, \tag{34}$$

$$w \in \mathcal{W}. \tag{35}$$

The linear program has an exponential number of constraints, but the nominal problem gives a separation oracle.

Theorem 2. *For discrete scenario uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized adversarial maxmin regret problem $\max_{w\in\mathcal{W}} \min_{T\in\mathcal{F}} \sum_{S\in\mathcal{S}} w_S \left(F(T,c^S) - F^*(c^S)\right)$ can be solved in time polynomial in n, m, and k.*

Proof. The separation oracle for (34) is given by the nominal problem. First, notice that $F^*(c^S)$ for $S \in \mathcal{S}$ can be computed once at initialization and then stored for easy computation of $\sum_{S\in\mathcal{S}} w_S F^*(c^S)$ for any w. Next, we have

$$\sum_{S\in\mathcal{S}} w_S F(T,c^S) = \sum_{S\in\mathcal{S}} w_S \sum_{e\in T} c_e^S = \sum_{e\in T} \left(\sum_{S\in\mathcal{S}} w_S c_e^S \right). \tag{36}$$

This means that solving nominal problem with costs $d = (d_1, \ldots, d_n)$ where

$$d_e = \sum_{S\in\mathcal{S}} w_S c_e^S \tag{37}$$

and comparing the solution with z and $\sum_{S\in\mathcal{S}} w_S F^*(c^S)$ gives the oracle. □

Corollary 2. *For discrete scenario uncertainty, $Z_{\mathrm{R}} = Z_{\mathrm{AR}}$.*

Proof. Using the substitution of the marginal probability vector in (23), it can be verified that the linear program solved by the adversary (33) - (35) is the dual of the program solved by the optimizing player (28) - (30). The result holds by strong duality. □

3.3 Primal-Dual Approximation

As noted in the above corollary, the linear program solved by the adversary (33) - (35) is the dual of program solved by the optimizing player (28) - (30). These linear programs correspond to the relaxation of the deterministic minmax regret problem, and can thus be used to develop a primal-dual approximation scheme. We will refer to the program solved by the optimizing player as the primal linear program, and the problem solved by the adversary as the dual linear program.

We rewrite the dual program (33) - (35) as

$$\max \quad z - \sum_{S\in\mathcal{S}} w_S F^*(c^S) \tag{38}$$

$$\text{s.t.} \quad \sum_{S\in\mathcal{S}} w_S F(T,c^S) \geq z, \qquad \forall T \in \mathcal{F}, \tag{39}$$

$$w \in \mathcal{W}. \tag{40}$$

A simple feasible solution to this program is given first by setting $w_S = 1/k$ for each $S \in \mathcal{S}$. Using the standard approach for primal-dual algorithms [26], we start with a sufficiently small value of z and increase it until a constraint becomes tight. The set corresponding to the tight solution is then added to the primal solution. The constraint (39) can be written as

$$\sum_{S \in \mathcal{S}} w_S F(T, c^S) = \sum_{S \in \mathcal{S}} w_S \sum_{e \in T} c_e^S = \sum_{e \in T} \left(\frac{1}{k} \sum_{S \in \mathcal{S}} c_e^S \right). \tag{41}$$

The first constraint that becomes tight corresponds to the set M that minimizes the mean costs over all scenarios,

$$M := \operatorname*{argmin}_{T \in \mathcal{F}} \sum_{e \in T} \left(\frac{1}{k} \sum_{S \in \mathcal{S}} c_e^S \right). \tag{42}$$

The set M, which is complete primal feasible solution, is added to the primal problem. Additionally, we have a feasible solution to the adversarial (dual) linear program with objective value

$$\left(\frac{1}{k} \right) \sum_{S \in \mathcal{S}} \left(\sum_{e \in M} c_e^S - F^*(c^S) \right), \tag{43}$$

which is a lower bound for the optimal objective value Z_{R}. Using the same observations made in [3], this gives a k-approximation algorithm for the minmax regret problem. The result given by the primal-dual framework is stronger than the result proved in [3] since it bounds the value of the approximate solution within a factor k of the linear programming relaxation value Z_{R}, rather than the integer program value Z_{D}.

Theorem 3. *For discrete scenario uncertainty, the solution to the nominal problem with mean costs is a k-approximation algorithm for the deterministic minmax regret problem.*

Proof. Using the construction above for a lower bound on Z_{R}, we have

$$\frac{Z_{\mathrm{D}}}{k} \leq \left(\frac{1}{k} \right) \max_{S \in \mathcal{S}} \left(\sum_{e \in M} c_e^S - F^*(c^S) \right) \leq \left(\frac{1}{k} \right) \sum_{S \in \mathcal{S}} \left(\sum_{e \in M} c_e^S - F^*(c^S) \right) \leq Z_{\mathrm{R}}. \tag{44}$$

The first inequality follows by definition of the deterministic minmax regret, the second inequality by a simple identity between the sum of a set of values and the maximum, and the third inequality from the linear program. □

An interesting corollary is a tight bound on the power of randomization in the minmax regret problem. For any nominal problem, moving from a deterministic solution to a randomized solution allows the optimizing player to at most reduce the expected regret by a factor of k.

Corollary 3. *For discrete scenario uncertainty,*

$$Z_{\mathrm{R}} \geq \frac{Z_{\mathrm{D}}}{k}. \tag{45}$$

The corollary equivalently states that the *integrality gap*, defined as the largest possible ratio of the optimal objective value of a program to its optimal linear programming relaxation, is equal to k. This holds independent of the nominal problem. We construct a tight example for the corollary using $n = k$ items, where the goal of the problem is simply to select the single item with lowest cost. For each item, there exists a scenario where the item has cost $c_e = 1$ and all other items have costs $c_e = 0$. The deterministic minmax regret is equal to 1 for the problem. In the randomized problem, the optimizing player selects each item with probability $1/k$ and the adversary selects each scenario with probability $1/k$. The expected regret is equal to the probability that the optimizing player selects the same item that the adversary assigns unit cost to, which is equal to $1/k$.

4 Interval Uncertainty

In this section we assume that cost uncertainty is characterized by interval uncertainty, meaning that each item cost is independently contained within known lower and upper bounds:

$$c_e \in [c_e^-, c_e^+], \quad \forall e \in E. \tag{46}$$

Define the region

$$\mathcal{I} := \{c | c_e \in [c_e^-, c_e^+], e \in E\}. \tag{47}$$

The set $\mathcal{I}$ is in general infinite. Since we wish to use a mixed distribution over $\mathcal{I}$ with finite support, we loosely define the set $\mathcal{I}_f$ to be some subset $\mathcal{I}_f \subset \mathcal{I}$ with finite cardinality, over which a probability distribution will be defined. The exact construction of $\mathcal{I}_f$ will become clear during the analysis, but a sufficient example is the set of cost vectors where costs are set equal to their lower or upper bounds, $\mathcal{I}_f = \{c | c_e = c_e^- \text{ or } c_e = c_e^+, e \in E\}$.

We proceed in the same way as the last section, studying the optimal policy for the optimizing player and then the adversary, followed by a primal-dual approximation algorithm for the deterministic problem. We restate notation and definitions throughout.

4.1 Optimizing Player

Under interval uncertainty, we have the deterministic maximum regret problem

$$R_{\max}(T) = \max_{c \in \mathcal{I}} R(T, c) = \max_{c \in \mathcal{I}} \left(F(T, c) - F^*(c) \right) \tag{48}$$

and the deterministic minmax regret problem

$$Z_{\mathrm{D}} = \min_{T \in \mathcal{F}} R_{\max}(T) = \min_{T \in \mathcal{F}} \max_{c \in \mathcal{I}} (F(T, c) - F^*(c)). \tag{49}$$

The deterministic minmax regret problem is well studied and can be solved with a mixed integer program [18]. We use an unconventional formulation, which has an exponential number of constraints. We will ultimately show that the randomized minmax regret problem corresponds to the linear programming relaxation of this formulation.

Lemma 4. *For interval uncertainty, the deterministic minmax regret problem* (6) *is equivalent to the following integer program.*

$$\begin{aligned} Z_{\mathrm{D}} \;=\; \min \quad & z & (50)\\ \text{s.t.} \quad & \sum_{e \in E \setminus T} c_e^+ x_e - \sum_{e \in T} c_e^- (1 - x_e) \le z, \qquad \forall T \in \mathcal{F}, \\ & x \in \mathcal{X}. \end{aligned}$$

Proof. From the maximum regret definition (48) and using vector notation instead of set notation,

$$\begin{aligned} R_{\max}(x) &= \max_{c\in\mathcal{I}} \left(F(x,c) - F^*(c)\right) \\ &= \max_{c\in\mathcal{I}} \left(\sum_{e\in E} c_e x_e - \min_{T\in\mathcal{F}} \sum_{e\in T} c_e\right) \\ &= \max_{T\in\mathcal{F}} \max_{c\in\mathcal{I}} \left(\sum_{e\in E} c_e x_e - \sum_{e\in T} c_e\right) \\ &= \max_{T\in\mathcal{F}} \max_{c\in\mathcal{I}} \left(\sum_{e\in E\setminus T} c_e x_e - \sum_{e\in T} c_e(1-x_e)\right) \\ &= \max_{T\in\mathcal{F}} \left(\sum_{e\in E\setminus T} c_e^+ x_e - \sum_{e\in T} c_e^-(1-x_e)\right), \end{aligned} \tag{51}$$

where in the third equality we have used that the expression $\sum_{e\in E} c_e x_e$ is not a function of T, and the last equality follows since $x_e \in \{0,1\}$. The program is then valid by the definition of the maximum. □

In the randomized model, the maximum expected regret is

$$\begin{aligned} \overline{R}_{\max}(y) &= \max_{c\in\mathcal{I}} \sum_{T\in\mathcal{F}} y_T R(T,c) \\ &= \max_{c\in\mathcal{I}} \left(\sum_{T\in\mathcal{F}} y_T F(T,c) - F^*(c)\right). \end{aligned} \tag{52}$$

As with the discrete scenario uncertainty case, we define the expected value of a solution for a distribution y and cost vector c,

$$\overline{F}(y,c) := \sum_{T\in\mathcal{F}} y_T F(T,c), \tag{53}$$

so the maximum expected regret can be stated as

$$\overline{R}_{\max}(y) = \max_{c\in\mathcal{I}} (\overline{F}(y,c) - F^*(c)). \tag{54}$$

The randomized minmax regret problem is thus

$$Z_{\mathrm{R}} = \min_{y\in\mathcal{Y}} \overline{R}_{\max}(y) = \min_{y\in\mathcal{Y}} \max_{c\in\mathcal{I}} (\overline{F}(y,c) - F^*(c)). \tag{55}$$

Starting with analysis of the maximum expected regret problem (52), we use the same substitution that we used in the previous section. Specifically, we let

$$p_e = \sum_{U\in\mathcal{F}: e\in U} y_U, \quad e \in E, \tag{56}$$

and define the marginal probability vector $p = (p_1, \ldots, p_n)$. Slightly abusing notation, we write $\overline{R}_{\max}(p)$ in place of $\overline{R}_{\max}(y)$ via this substitution.

Lemma 5. *For interval uncertainty, the maximum expected regret problem* (52) *is equivalent to the problem*

$$\overline{R}_{\max}(p) = \max_{T\in\mathcal{F}} \left(\sum_{e\in E\setminus T} c_e^+ p_e - \sum_{e\in T} c_e^- (1-p_e) \right). \tag{57}$$

Proof. We start with (52) and use the substitution of p. The analysis is nearly identical to the proof of Lemma 4.

$$\begin{aligned}
\overline{R}_{\max}(y) &= \max_{c\in\mathcal{I}} (\overline{F}(y,c) - F^*(c)) \\
&= \max_{c\in\mathcal{I}} \left(\sum_{U\in\mathcal{F}} y_U \sum_{e\in U} c_e - \min_{T\in\mathcal{F}} \left(\sum_{e\in T} c_e \right) \right) \\
&= \max_{c\in\mathcal{I}} \left(\sum_{e\in E} c_e \sum_{U\in\mathcal{F}:e\in U} y_U - \min_{T\in\mathcal{F}} \left(\sum_{e\in T} c_e \right) \right) \\
&= \max_{c\in\mathcal{I}} \left(\sum_{e\in E} c_e p_e - \min_{T\in\mathcal{F}} \left(\sum_{e\in T} c_e \right) \right).
\end{aligned} \tag{58}$$

Now using the notation $\overline{R}_{\max}(p)$,

$$\begin{aligned}
\overline{R}_{\max}(p) &= \max_{c\in\mathcal{I}} \left(\sum_{e\in E} c_e p_e - \min_{T\in\mathcal{F}} \left(\sum_{e\in T} c_e \right) \right) \\
&= \max_{c\in\mathcal{I}} \max_{T\in\mathcal{F}} \left(\sum_{e\in E} c_e p_e - \sum_{e\in T} c_e \right) \\
&= \max_{T\in\mathcal{F}} \max_{c\in\mathcal{I}} \left(\sum_{e\in E\setminus T} c_e p_e - \sum_{e\in T} c_e (1-p_e) \right),
\end{aligned} \tag{59}$$

where in the first equality we have used that the expression $\sum_{e\in E} c_e p_e$ is not a function of T, and the other equalities follow from rearranging terms. Notice in (59) that p_e is simply the total probability that item e is selected, so for $y \in \mathcal{Y}$, we must have $p_e \in [0,1]$. This makes it easy to see that for a given $T \in \mathcal{F}$,

$$\max_{c\in\mathcal{I}} \left(\sum_{e\in E\setminus T} c_e p_e - \sum_{e\in T} c_e (1-p_e) \right) = \sum_{e\in E\setminus T} c_e^+ p_e - \sum_{e\in T} c_e^- (1-p_e). \tag{60}$$

Substituting (60) into (59) then gives an optimization problem with a finite number of feasible solutions,

$$\overline{R}_{\max}(p) = \max_{T\in\mathcal{F}} \left(\sum_{e\in E\setminus T} c_e^+ p_e - \sum_{e\in T} c_e^- (1-p_e) \right), \tag{61}$$

which completes the proof. □

An immediate corollary of Lemma 5 is that we can solve the maximum expected regret problem for a given y by enumerating all $|\mathcal{F}|$ subsets (potentially an exponential number of them) and choosing the one that maximizes the argument of (61). This allows the entire randomized minmax regret problem to be restated.

Lemma 6. *For interval uncertainty, the objective value Z_{R} of the randomized minmax regret problem* (11) *is equal to that of the problem*

$$\min_{p \in \mathrm{CH}(\mathcal{X})} \max_{T \in \mathcal{F}} \left(\sum_{e \in E \setminus T} c_e^+ p_e - \sum_{e \in T} c_e^- (1 - p_e) \right), \tag{62}$$

where $\mathrm{CH}(\mathcal{X})$ *denotes the convex hull of* $\mathcal{X}$.

Proof. By the same argument as the proof of Lemma 2. □

Using Lemma 6, we can now formulate a linear program to solve the randomized minmax regret problem:

$$\begin{aligned}
\min \quad & z && (63)\\
\text{s.t.} \quad & \sum_{e \in E \setminus T} c_e^+ p_e - \sum_{e \in T} c_e^- (1 - p_e) \le z, \qquad \forall T \in \mathcal{F}, && (64)\\
& p \in \mathrm{CH}(\mathcal{X}). && (65)
\end{aligned}$$

While the above program may have an exponential number of constraints, it can be solved efficiently via the ellipsoid algorithm if a separation oracle is available for the constraints (64) and (65). This brings us to our main result.

Theorem 4. *For interval uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized minmax regret problem* $\min_{y \in \mathcal{Y}} \max_{c \in \mathcal{I}} (\overline{F}(y, c) - F^*(c))$ *can be solved in time polynomial in n and m.*

Proof. Consider the linear program (63) - (65). The separation oracle for the constraints (65) is given by the equivalence of optimization and separation [16]. To see the separation oracle for the constraint (64), we define the item cost vector $d = (d_1, \ldots, d_n)$ where

$$d_e = c_e^- + p_e(c_e^+ - c_e^-), \quad e \in E, \tag{66}$$

and then solve

$$z_d = \min_{T \in \mathcal{F}} \sum_{e \in T} d_e. \tag{67}$$

Let T_d be the set that minimizes the above expression. If $\sum_{e \in E} c_e^+ p_e - z_d \le z$, then we are guaranteed feasibility, otherwise the separating hyperplane (64) is generated where $T = T_d$. To see

the validity of this approach, we have

$$\begin{aligned}
\sum_{e\in E} c_e^+ p_e - z_d &= \sum_{e\in E} c_e^+ p_e - \min_{T\in\mathcal{F}} \sum_{e\in T} d_e \\
&= \sum_{e\in E} c_e^+ p_e - \min_{T\in\mathcal{F}} \sum_{e\in T} (c_e^- + p_e(c_e^+ - c_e^-)) \\
&= \max_{T\in\mathcal{F}} \left(\sum_{e\in E} c_e^+ p_e - \sum_{e\in T} (c_e^- + p_e(c_e^+ - c_e^-)) \right) \\
&= \max_{T\in\mathcal{F}} \left(\sum_{e\in E\setminus T} c_e^+ p_e - \sum_{e\in T} c_e^- (1-p_e) \right). \qquad (68)
\end{aligned}$$

The solution to the linear program (63) - (65) is a vector p, which can then be used to find a mixed strategy y in polynomial time using Lemma 3. □

Corollary 4. *For interval uncertainty,* $Z_{\mathrm{R}} \leq Z_{\mathrm{D}}$.

Proof. This follows simply by noting that the program (63) - (65) is the linear programming relaxation of (50). □

4.2 Adversary

The set $\mathcal{I}_f$ is necessary for describing the distribution of the adversary. The distribution over costs is $w = (w_c)_{c\in\mathcal{I}_f}$ and $\mathcal{W}$ again indicates the set of valid distributions. The minimum expected regret problem is

$$\begin{aligned}
\overline{R}_{\min}(w) &= \min_{T\in\mathcal{F}} \sum_{c\in\mathcal{I}_f} w_c R(T,c) \\
&= \min_{T\in\mathcal{F}} \sum_{c\in\mathcal{I}_f} w_c \left(F(T,c) - F^*(c) \right). \qquad (69)
\end{aligned}$$

The adversarial randomized maxmin regret problem is then

$$Z_{\mathrm{AR}} = \max_{w\in\mathcal{W}} \overline{R}_{\min}(w) = \max_{w\in\mathcal{W}} \min_{T\in\mathcal{F}} \sum_{c\in\mathcal{I}_f} w_c \left(F(T,c) - F^*(c) \right). \qquad (70)$$

We can directly formulate a linear program for the adversarial maxmin regret problem, explicitly writing the constraints for $w \in \mathcal{W}$.

$$\begin{aligned}
\max \quad & z && (71)\\
\text{s.t.} \quad & \sum_{c\in\mathcal{I}_f} w_c \left(\sum_{e\in T} c_e - F^*(c) \right) \geq z, \qquad \forall T \in \mathcal{F}, && (72)\\
& \sum_{c\in\mathcal{I}_f} w_c = 1, && (73)\\
& w_c \geq \mathbf{0}. && (74)
\end{aligned}$$

Since this program has an exponential number of constraints and potentially an exponential number of variables, we consider its dual. We expect this "dual of the dual" program to be the primal linear program solved by the optimizing player; this will indeed be the case after some manipulation. The dual of (71) - (74) is

$$\min \quad \beta \tag{75}$$

$$\text{s.t.} \quad \sum_{T\in\mathcal{F}} \alpha_T \left(\sum_{e\in T} c_e - F^*(c) \right) \leq \beta, \qquad \forall c \in \mathcal{I}_f, \tag{76}$$

$$\sum_{T\in\mathcal{F}} \alpha_T = 1, \tag{77}$$

$$\alpha_T \geq \mathbf{0}. \tag{78}$$

To simplify, note that for a feasible $\alpha = (\alpha_T)_{T\in\mathcal{F}}$, we have $\sum_{T\in\mathcal{F}} \alpha_T F^*(c) = F^*(c)$. Furthermore,

$$\sum_{T\in\mathcal{F}} \alpha_T \sum_{e\in T} c_e = \sum_{e\in E} c_e \sum_{T\in\mathcal{F}:e\in T} \alpha_T. \tag{79}$$

We use the substitution

$$q_e := \sum_{T\in\mathcal{F}:e\in T} \alpha_T, \quad e \in E. \tag{80}$$

Let $q = (q_1, \ldots, q_n)$. The substitution yields the linear program

$$\min \quad \beta \tag{81}$$

$$\text{s.t.} \quad \sum_{e\in E} c_e q_e - F^*(c) \leq \beta, \qquad \forall c \in \mathcal{I}_f, \tag{82}$$

$$q \in \text{CH}(\mathcal{X}). \tag{83}$$

This program no longer has an exponential number of variables, and the exponential number of constraints can be handled via separation, which we describe shortly. First, for some set $A \in \mathcal{F}$, define the cost vector $c^A = (c^A_e)_{e\in E}$ where

$$c^A_e := \begin{cases} c^-_e, & e \in A, \\ c^+_e, & e \in E \setminus A. \end{cases} \tag{84}$$

That is, c^A is the cost vector where all costs are equal to their upper bound, except for costs in the set A, which are equal to their lower bound. The theorem below shows that without loss of generality, we can can consider cost vectors of this form for separation of the constraint (82). This allows us to define $\mathcal{I}_f$ as the set of all cost vectors $\{c^A,\ A \in \mathcal{F}\}$. Since this may still be an exponentially sized set, we define an adversarial mixed strategy encoding $\mathcal{L} = (C, W)$ as a set of costs $C = \{c^{A_j} \in \mathcal{I}_f,\ j = 1, \ldots, \eta\}$ to be selected with corresponding probabilities $W = \{w_{c^{A_j}} \in [0, 1],\ j = 1, \ldots, \eta\}$ satisfying $\sum_{j=1}^{\eta} w_{c^{A_j}} = 1$.

Theorem 5. *For interval uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized adversarial maxmin regret problem $\max_{w\in\mathcal{W}} \min_{T\in\mathcal{F}} \sum_{c\in\mathcal{I}_f} w_S \left(F(T, c) - F^*(c)\right)$ can be solved in time polynomial in n and m.*

Proof. The constraint (82) is simply the maximum regret problem for a given vector q and can be generated via the nominal problem,

$$\max_{c\in\mathcal{I}_f}\left(\sum_{e\in E} c_e q_e - F^*(c)\right) \;=\; \max_{T\in\mathcal{F}}\left(\sum_{e\in E\setminus T} c_e^+ q_e - \sum_{e\in T} c_e^-(1-q_e)\right), \tag{85}$$

where we have used the analysis in Lemma 5. This allows us to write the linear program as

$$\min \quad \beta \tag{86}$$

$$\text{s.t.} \quad \sum_{e\in E\setminus T} c_e^+ q_e - \sum_{e\in T} c_e^-(1-q_e) \le \beta, \qquad \forall T\in\mathcal{F}, \tag{87}$$

$$q \in \mathrm{CH}(\mathcal{X}). \tag{88}$$

Note this is precisely the linear program (63) - (65) solved by the optimizing player. This justifies the assumption of the finite set $\mathcal{I}_f$: only a polynomial number of separating cost vectors will be generated, and they will be of the form c^A as defined in (84). The adversary is of course interested in the dual variables of the linear program (86) - (88). Each separating hyperplane generated for the constraint (87) gives a set $T\in\mathcal{F}$ for which the adversary adds the cost vector c^T to his mixed strategy; this cost vector has probability w_{c^T} in the mixed strategy, given by the corresponding dual variable. □

Corollary 5. *For interval uncertainty,* $Z_{\mathrm{R}} = Z_{\mathrm{AR}}$.

Proof. By strong duality, since the optimizing player solves (63)-(65), which is the dual of the linear program (71) - (74) for the adversary. □

4.3 Primal-Dual Approximation

The primal linear program for the optimizing player is (63)-(65) and the dual linear program for the adversary is (71) - (74). Using a similar approach to the previous section, we devise a primal-dual approximation algorithm for the deterministic minmax regret problem.

We rewrite the dual linear program as

$$\max \quad z - \sum_{c\in\mathcal{I}_f} w_c F^*(c) \tag{89}$$

$$\text{s.t.} \quad \sum_{c\in\mathcal{I}_f} w_c \sum_{e\in U} c_e \ge z, \qquad \forall U\in\mathcal{F}, \tag{90}$$

$$\sum_{c\in\mathcal{I}_f} w_c = 1, \tag{91}$$

$$w_c \ge \mathbf{0}. \tag{92}$$

We must select a feasible solution for w; we will make the simple choice of setting $w_c = 1/2$ for two cost vectors. Recall the definition of the cost vector c^A. For some set $A\in\mathcal{F}$, we have $c^A = (c_e^A)_{e\in E}$ where

$$c_e^A = \begin{cases} c_e^-, & e\in A, \\ c_e^+, & e\in E\setminus A, \end{cases} \tag{93}$$

Additionally, we define $c^{\overline{A}} = (c_e^{\overline{A}})_{e \in E}$ where

$$c_e^{\overline{A}} = \begin{cases} c_e^+, & e \in A, \\ c_e^-, & e \in E \setminus A. \end{cases} \tag{94}$$

We set $w_c = 1/2$ for $c = c^A$ and $c = c^{\overline{A}}$, so the linear program becomes

$$\max \quad z - \frac{1}{2}\left(F^*(c^A) - F^*(c^{\overline{A}})\right) \tag{95}$$

$$\text{s.t.} \quad \sum_{e \in U} \left(\frac{c_e^- + c_e^+}{2}\right) \geq z, \qquad \forall U \in \mathcal{F}. \tag{96}$$

Under the primal dual approach, we increase z until one of the constraints becomes tight. The first tight constraint corresponds to the primal solution M that has minimum total cost under the midpoint costs:

$$M := \operatorname*{argmin}_{U \in \mathcal{F}} \sum_{e \in U} \left(\frac{c_e^- + c_e^+}{2}\right). \tag{97}$$

This gives the objective value

$$\frac{1}{2}\left(\sum_{e \in M}(c_e^- + c_e^+) - F^*(c^A) - F^*(c^{\overline{A}})\right). \tag{98}$$

By choosing the set A to be the midpoint cost minimizing set M, we can write the resulting objective value in terms of the maximum deterministic regret.

Lemma 7. *For $A = M$,*

$$\sum_{e \in M}(c_e^- + c_e^+) - F^*(c^A) - F^*(c^{\overline{A}}) = R_{\max}(M). \tag{99}$$

Proof. The maximum regret for the set M can be expressed as

$$R_{\max}(M) = \sum_{e \in M} c_e^+ - \min_{T \in \mathcal{F}} \left(\sum_{e \in T \cap M} c_e^+ + \sum_{e \in T \setminus M} c_e^-\right). \tag{100}$$

Also note that

$$F^*(c^{\overline{A}}) = \min_{T \in \mathcal{F}} \left(\sum_{e \in T \cap A} c_e^+ + \sum_{e \in T \setminus A} c_e^-\right). \tag{101}$$

Thus for $A = M$, we have

$$R_{\max}(M) = \sum_{e \in M} c_e^+ - F^*(c^{\overline{A}}). \tag{102}$$

It is left to show that $F^*(c^M) = \sum_{e \in M} c_e^-$. This, however, immediately follows with a simple argument. Since the set M is minimum for midpoint costs, it must also be minimum for costs c^M (i.e., the costs where $c_e = c_e^-$ for all $e \in M$ and $c_e = c_e^+$ for all $e \in E \setminus M$). □

This gives a new proof that solving the nominal problem with midpoint costs gives a 2-approximation to the deterministic minmax regret problem. Once again, the result here is stronger than the result of [19] since it states that the value of the approximate solution is within a factor 2 of Z_{R} rather than just Z_{D}.

Theorem 6. *For interval uncertainty, the solution to the nominal problem with midpoint costs is a 2-approximation algorithm for the deterministic minmax regret problem.*

Proof. The linear program listed above is the dual of the problem solved by the optimizing player in the randomized framework. By weak duality, any feasible solution to the above program gives a lower bound on the value of the game in the randomized framework, Z_{R}. The construction described above using costs c^A and $c^{\overline{A}}$ gives a feasible solution to the program, and Lemma 7 allows us to express the resulting objective value in terms of the maximum deterministic regret for a solution set M using $A = M$. Specifically,

$$\frac{Z_{\mathrm{D}}}{2} \quad \leq \quad \frac{R_{\max}(M)}{2} \leq Z_{\mathrm{R}} \leq Z_{\mathrm{D}}, \tag{103}$$

where the first inequality follows from the definition of the deterministic minmax regret problem, the second inequality follows using Lemma 7 with the feasible linear programming solution, and the third inequality follows from Theorem 4. □

The potential gain from using randomization under interval uncertainty is not as significant as with discrete scenario uncertainty.

Corollary 6. *For interval uncertainty,*

$$Z_{\mathrm{R}} \geq \frac{Z_{\mathrm{D}}}{2}. \tag{104}$$

Independently of the nominal problem, the integrality gap for the minmax regret problem is equal to 2. A tight example for the corollary is easily constructed. Consider a problem with two items $E = \{e_1, e_2\}$ where the optimizing player must choose one item. Let $(c_e^+, c_e^-) = (0, 1)$ for both items $e = e_1, e_2$. It can be verified that $Z_{\mathrm{D}} = 1$ and $Z_{\mathrm{R}} = 1/2$.

5 General Uncertainty Sets

In this section, we show that if the uncertainty set $\mathcal{C}$ is allowed to be a general nonnegative convex set and the nominal problem is polynomial solvable, the maximum expected regret problem becomes NP-hard. Note that the deterministic maximum regret problem is a special case of the maximum expected regret problem. The result of this section thus implies that both randomized and deterministic minmax regret problems are NP-hard under general convex uncertainty, even if the nominal problem is polynomial solvable.

We restate the maximum expected regret problem for general uncertainty sets. For a given marginal probability vector p, the maximum expected regret problem is, starting with the first line

of (59),

$$\begin{aligned}\overline{R}_{\max}(p) &= \max_{c\in\mathcal{C}} \left(\sum_{e\in E} c_e p_e - F^*(c) \right) \\ &= \max_{c\in\mathcal{C}} \left(\sum_{e\in E} c_e p_e - \min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e \right) \\ &= \max_{c\in\mathcal{C}} \max_{x\in\mathcal{X}} \left(\sum_{e\in E} c_e (p_e - x_e) \right). \end{aligned} \tag{105}$$

Negating the objective function, the maximum expected regret problem is equivalent to

$$-\overline{R}_{\max}(p) = \min_{c\in\mathcal{C}} \min_{x\in\mathcal{X}} \sum_{e\in E} c_e (x_e - p_e), \tag{106}$$

for a given $p \in \mathrm{CH}(\mathcal{X})$.

Before we examine the complexity of (106), we consider the following problem, which we refer to as the *bilinear combinatorial problem*:

$$\min_{c\in\mathcal{C}} \min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e. \tag{107}$$

We demonstrate the hardness of this problem via a reduction from the Hamiltonian path problem; the proof is similar to the standard proof for showing that the intersection of three matroids is NP-hard.

Lemma 8. *For polynomial solvable nominal problems $F^*(c) = \min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e$ and nonnegative convex uncertainty sets $\mathcal{C}$, the bilinear combinatorial problem $\min_{c\in\mathcal{C}} \min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e$ is NP-hard.*

Proof. Recall that the directed Hamiltonian path problem asks the following: given a directed graph $G = (V, E)$ with a designated source node s and terminal node t, does there exist a path starting at s and ending at t that visits each node exactly once? For a given instance of the directed Hamiltonian path problem, we construct an instance of the bilinear combinatorial problem such that it has an optimal objective value of zero if an only if the graph contains a valid Hamiltonian path.

We construct the set $\mathcal{C}$ to indicate the selection of edges such that each vertex has exactly one incoming edge (except for vertex s) and one outgoing edge (except for vertex t). Specifically, we say that an edge e is selected if its cost c_e is equal to zero, otherwise we refer to it as blocked. The notation $\delta^+(v)$ (respectively $\delta^-(v)$) indicates the set of outgoing (incoming) edges for vertex

v. The constraints for the set $\mathcal{C}$ are

$$\begin{aligned}
\sum_{e\in\delta^-(v)} c_e &= |\delta^-(v)| - 1, \quad v \in V \setminus \{s\}, \\
\sum_{e\in\delta^-(s)} c_e &= |\delta^-(s)|, \\
\sum_{e\in\delta^+(v)} c_e &= |\delta^+(v)| - 1, \quad v \in V \setminus \{t\}, \\
\sum_{e\in\delta^+(t)} c_e &= |\delta^+(t)|, \\
0 \leq c_e \leq 1, &\qquad e \in E.
\end{aligned} \tag{108}$$

Note that for a given vertex, if one if its incoming edges is selected ($c_e = 0$), then all the remaining incoming edges must be blocked ($c_e = 1$); the same holds for outgoing edges.

Define $\mathcal{X}$ to indicate the set of all feasible spanning trees for G, so that $\min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e$ is the minimum spanning tree problem. For an optimal solution $x^* = (x_1^*, \ldots, x_n^*)$ to the bilinear combinatorial problem giving zero objective value, the set $\{e | x_e^* = 0, e \in E\}$ indicates a valid Hamiltonian path, as the construction of $\mathcal{C}$ indicates that all vertices have one selected incoming and outgoing edge (except for s and t), and the spanning tree ensures that no cycles are present. Finally, we have that the construction of the set $\mathcal{C}$ can be done in polynomial time. □

Theorem 7. *For polynomial solvable nominal problems* $F^*(c) = \min_{x\in\mathcal{X}} \sum_{e\in E} c_e x_e$ *and nonnegative convex uncertainty sets* $\mathcal{C}$*, the maximum expected regret problem* $\max_{c\in\mathcal{C}} \left(\sum_{e\in E} c_e p_e - F^*(c)\right)$ *where* $p \in \mathrm{CH}(\mathcal{X})$ *is* NP*-hard.*

Proof. We modify the reduction used for Lemma 8 to account for the presence of some $p \in \mathrm{CH}(\mathcal{X})$ in the objective function of the maximum expected regret problem, which is now

$$\min_{c\in\mathcal{C}} \min_{x\in\mathcal{X}} \sum_{e\in E} c_e (x_e - p_e). \tag{109}$$

Again let $\mathcal{X}$ indicate the set of all feasible spanning trees for the directed graph $G = (V, E)$ and let $p \in \mathcal{X}$ be a valid spanning tree. We construct a new graph over the same set of vertices by taking G and duplicating $|V| - 1$ edges. For each edge given by the spanning tree p, we choose a corresponding edge in G (note that there may be more than one option if both edges (v_i, v_j) and (v_j, v_i) are present, for example) and duplicate it. Let this new graph be denoted by $G' = (V, E')$, and let the set of all spanning trees over the new graph be indicated by $\mathcal{X}'$. Let $\widetilde{p} \in \mathcal{X}'$ indicate the set of edges that were constructed via duplication (i.e. the edges $E' \setminus E$), which is a valid spanning tree for G'. We finally construct the set $\mathcal{C}'$ using the inequalities in (108) but over E' instead of E.

Now consider the modified maximum expected regret problem

$$\min_{c\in\mathcal{C}'} \min_{x\in\mathcal{X}'} \sum_{e\in E'} c_e (x_e - \widetilde{p}_e) = \min_{c\in\mathcal{C}'} \min_{x\in\mathcal{X}'} \left(\sum_{e\in E} c_e x_e + \sum_{e\in E'\setminus E} c_e (x_e - 1) \right). \tag{110}$$

It can be seen that the modified problem has an objective value equal to $-(|V| - 1)$ if and only if G has a Hamiltonian path. This corresponds to the first sum in objective function being equal

to zero, and the second sum being equal to $-(|V|-1)$. As before, for an optimal solution $x^* = (x_1^*, \ldots, x_{n+|V|-1}^*)$ to the modified problem, the set $\{e|x_e^* = 0, e \in E'\}$ gives a Hamiltonian path that is valid for both G' and G. Notice that all of the duplicated edges $e \in E' \setminus E$ must be blocked ($c_e = 1$) and not selected by the minimum spanning tree ($x_e = 0$) for the objective value to be equal to $-(|V|-1)$. To finish the proof, we observe that the construction of $\mathcal{C}'$ and G' can be accomplished in polynomial time. □

6 Conclusion

We have shown that for both the interval and discrete scenario representations of uncertainty, the randomized minmax regret version of any polynomial solvable combinatorial problem is polynomial solvable. Furthermore, the maximum expected regret in the randomized model is upper bounded by the maximum regret of the deterministic model. These results, including the fact that there always exists a polynomial-sized optimal solution for randomized minmax regret, are at first glance somewhat surprising. Intuitively, the polynomial solvability of the randomized model results from the fact that a linear program must be solved instead of the integer program or mixed integer program (which is required for the deterministic model). The improvement in performance holds because the adversary has less power in the randomized model than the deterministic model.

For many applications that are not adversarial in nature, the randomized minmax regret criteria is likely a more appropriate model than the deterministic version. In particular, the deterministic solution may be overly conservative since costs are not truly chosen in an adversarial fashion in response to the selected solution. On the other hand, one must be willing to tolerate higher variance if randomization is used.

Our results on lower bounds for randomized minmax regret in relation to deterministic minmax regret, specifically Corollary 3 and 6, have important implications for approximating deterministic minmax regret problems. In Kasperski [18], it is posed as an open problem whether or not there exist approximation algorithms for interval uncertainty that, for some specific nominal problems, achieve an approximation ratio better than 2. In some sense, we have answered this question in the negative. Corollary 6 indicates that the integrality gap for the minmax regret problem is equal to 2, and it is easy to create instances of nearly all nominal problems that achieve this gap. The same can be argued for the integrality gap of k under discrete scenario uncertainty. Integrality gaps bound the best possible performance that can be obtained from approximation algorithms based on linear programming relaxations, which includes many approximation techniques [26]. Nonetheless, an important shortcoming of our primal-dual algorithms is that they do not use optimal dual solutions. It may be possible to design algorithms that perform better on specific problem instances by using optimal dual solutions. Under discrete scenario uncertainty, for instance, solving the nominal problem with costs averaged by dual variable weights, rather than averaged uniformly, may give improved performance.

An important future step with randomized minmax regret research is to develop approximation algorithms (now in the randomized model) for dealing with nominal problems that are already NP-complete. This problem is non-trivial: an algorithm with an approximation factor α for a nominal problem does not immediately yield an algorithm to approximate the randomized minmax regret problem with a factor α. Another interesting topic to study from an experimental perspective is a hybrid approach that employs both deterministic and randomized minmax regret. For example, one could find a solution that minimizes the maximum expected regret, subject to the maximum regret being no greater than some constant. The algorithm for this problem can be easily constructed by

combining our results with existing work, but may no longer be polynomial solvable.

References

[1] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Complexity of the min–max and min–max regret assignment problems. *Operations Research Letters*, 33(6):634–640, 2005.

[2] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Complexity of the min-max (regret) versions of cut problems. In *Algorithms and Computation*, pages 789–798. Springer, 2005.

[3] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Approximating min-max (regret) versions of some polynomial problems. In *Computing and Combinatorics*, pages 428–438. Springer, 2006.

[4] Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Min–max and min–max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research*, 197(2):427–438, 2009.

[5] Igor Averbakh. On the complexity of a class of combinatorial optimization problems with uncertainty. *Mathematical Programming*, 90(2):263–272, 2001.

[6] Igor Averbakh and Vasilij Lebedev. Interval data minmax regret network optimization problems. *Discrete Applied Mathematics*, 138(3):289–301, 2004.

[7] Igor Averbakh and Vasilij Lebedev. On the complexity of minmax regret linear programming. *European Journal of Operational Research*, 160(1):227–231, 2005.

[8] Shai Ben-David, Allan Borodin, Richard Karp, Gabor Tardos, and Avi Wigderson. On the power of randomization in on-line algorithms. *Algorithmica*, 11(1):2–14, 1994.

[9] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1): 35–53, 2004.

[10] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization under ellipsoidal uncertainty sets. *Technical report, Massachusetts Institute of Technology*, 2004.

[11] Dimitris Bertsimas, Ebrahim Nasrabadi, and James B. Orlin. On the power of nature in robust discrete optimization. *In preparation*, 2012.

[12] Dimitris Bertsimas, Ebrahim Nasrabadi, and James B. Orlin. On the power of randomization in network interdiction. *In preparation*, 2012.

[13] John R Birge and François V Louveaux. *Introduction to stochastic programming*. Springer, 1997.

[14] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*, volume 53. Cambridge University Press Cambridge, 1998.

[15] Constantin Carathéodory. Über den variabilitätsbereich der fourierschen konstanten von positiven harmonischen funktionen. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 32(1):193–217, 1911.

[16] Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

[17] Berit Johannes and James B. Orlin. Minimax regret problems are harder than minimax problems. *Submitted*, 2012.

[18] Adam Kasperski. *Discrete optimization with interval data: minmax regret and fuzzy approach*, volume 228. Springer, 2008.

[19] Adam Kasperski and Paweł Zieliński. An approximation algorithm for interval data minmax regret combinatorial optimization problems. *Information Processing Letters*, 97(5):177–180, 2006.

[20] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *AAAI*, 2010.

[21] Panos Kouvelis and Gang Yu. *Robust discrete optimization and its applications*, volume 14. Springer, 1997.

[22] J v Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[23] Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.

[24] Alexander Shapiro, Darinka Dentcheva, and Andrzej P Ruszczyński. *Lectures on stochastic programming: modeling and theory*, volume 9. SIAM, 2009.

[25] Abraham Wald. Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics*, 10(4):299–326, 1939.

[26] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge University Press, 2011.

[27] Gang Yu and Jian Yang. On the robust shortest path problem. *Computers & Operations Research*, 25(6):457–468, 1998.